arXiv:1404.4744v1 [cs.CR] 18 Apr 2014

PrivLoc: Preventing Location Tracking in Geofencing Services

Jens Mathias Bohli, Dan Dobre, Ghassan O. Karame, Wenting Li

NEC Laboratories Europe, Germany
firstname.lastname@neclab.eu

Abstract. Location-based services are increasingly used in our daily activities. In current services, users however have to give up their location privacy in order to acquire the service.
The literature features a large number of contributions which aim at enhancing user privacy in location-based services. Most of these contributions obfuscate the locations of users using spatial and/or temporal cloaking in order to provide k-anonymity. Although such schemes can indeed strengthen the location privacy of users, they often decrease the service quality and do not necessarily prevent the possible tracking of user movements (i.e., direction, trajectory, velocity). With the rise of Geofencing applications, tracking of movements becomes more evident since, in these settings, the service provider is not only requesting a single location of the user, but requires the movement vectors of users to determine whether the user has entered/exited a Geofence of interest.
In this paper, we propose a novel solution, PrivLoc, which enables the privacy-preserving outsourcing of Geofencing and location-based services to the cloud without leaking any meaningful information about the location, trajectory, and velocity of the users. Notably, PrivLoc enables an efficient and privacy-preserving intersection of movement vectors with any polygon of interest, leveraging functionality from existing Geofencing services or spatial databases. We analyze the security and privacy provisions of PrivLoc and we evaluate the performance of our scheme by means of implementation. Our results show that the performance overhead introduced by PrivLoc can be largely tolerated in realistic deployment settings.

1 Introduction

Location-based services (e.g., Foursquare [1] and Yelp [5]) are gaining increasing importance recently. Several applications enable users (e.g., using mobile devices) to discover and communicate their locations to a server in the cloud; in turn, the server uses this information to return data relevant at the users' locations. For instance, a number of existing services can only be acquired by users who are located within a specific geographical area; these include banking services, Youtube, and content delivery services, among many others. Location information also proves to be useful for a number of security-critical services such as police investigations, e-voting, etc.

However, while many devices (e.g., smartphones, tablets) are capable of discovering and reporting their locations, a considerable number of users shy away from reporting their locations in the fear of being tracked or profiled by service providers [7]. This problem is even more evident when the service provider wishes to outsource his spatial services to the cloud (e.g., [4]). Service providers have considerable incentives to rely on hosted services in the cloud, since this enables them to maximize the availability of the service while minimizing the costs for acquisition of hardware and operation. Indeed, the cloud offers a low barrier for small and medium enterprises to offer location-based services and enables its clients to avoid huge upfront investments to accommodate for peak usage. However, hosting the spatial service in the cloud raises serious privacy concerns with respect to the leakage of client location information to the cloud provider.

The literature comprises a plethora of contributions which strengthen the privacy of users in location-based services. Most of these contributions focus on anonymizing user locations by means of a trusted location anonymizer server [6, 15, 20, 22, 34]. After registering with the service, users can send their exact locations to the server, which "blurs" these location reports and sends the cloaked location to a remote database server. The server also filters the database's response and subsequently sends the exact answers back to the users when needed. Existing location anonymization techniques can be categorized according to three different approaches: *(i)* inserting false dummies [21], where the server sends n location updates for each location reported by the user ($(n-1)$ reports of which are dummy), *(ii)* location blurring [6, 10, 15, 20, 33] where the location of the user is blurred into a spatial area (using spatial or temporal cloaking), and *(iii)* landmark obfuscation [19] where the server sends the location of a near-by landmark instead of the location of the user.

While these techniques can provide users with k-anonymity guarantees, existing techniques *(i)* often require changes to the database server in order to efficiently process the anonymized queries, or *(ii)* reduce the accuracy of the location-based service (e.g., when relying on spatial cloaking), or *(iii)* do not prevent location tracking [22]. Location tracking could be performed e.g., by inferring the direction of the movement/path followed by the user, the velocity of the user, etc., which might de-anonymize users. Such information leakage is particularity damaging in location-based services where service providers are interested in the events of users entering/exiting a given area (e.g., Geofencing applications [2, 3, 23, 28, 29]). These applications are gaining increasing importance for e.g., targeted advertisements, and typically take as inputs vectors of movements performed by users and enable service providers to extract various statistics about their customers, such as visit durations, start, end, etc.

In this paper, we address this problem, and we propose a novel solution, PrivLoc, which efficiently enables the privacy-preserving use of Geofencing services in the cloud without incurring any modifications to spatial indexing techniques, and without leaking any meaningful information about the location, trajectory, and/or velocity of the users to the cloud provider. More specifically, we

consider a typical Geofencing setting, whereby a spatial database hosted on the cloud receives regular movement vectors from users, and checks if these movements cross a Geofenced area which has been subscribed to within the database. If so, the provider issues a notification informing the appropriate subscriber that a user has exited/entered the subscribed area of interest. Note that this can be achieved in existing spatial databases by querying for intersection between the user movement vector and the Geofenced area. PrivLoc leverages the presence of a private trusted service which performs efficient specially-crafted transformations of location reports to interface with the cloud-hosted Geofencing service without exposing the privacy of users. As such, PrivLoc can be used by companies and individuals to prevent information leakage towards spatial databases hosted in clouds, such as Google, Amazon, etc. In this respect, PrivLoc enables a privacy-preserving intersection of movement vectors with any polygon of interest, while leveraging functionality from existing cloud-based spatial databases. We analyze the security and privacy provisions of PrivLoc and we evaluate its performance in a realistic setting. Our results show that our scheme scales well with the number of users and subscriptions in the system and does not incur considerable computational overhead on the trusted server.

The remainder of the paper is organized as follows. In Section 2, we outline our system and adversarial model. In Section 3, we introduce and analyze our solution, PrivLoc, which efficiently enables users to acquire privacy-preserving Geofencing services hosted in the cloud. We evaluate its performance by means of an implementation in Section 4. In Section 5, we overview related work in the area, and we conclude the paper in Section 6.

2 Model

In this section, we describe our system and adversarial model, and we outline the security requirements that our solution should satisfy.

2.1 Spatial Databases

Spatial databases are instances of databases optimized to store and query data which represents objects defined in a geometric space. Examples of spatial databases include MongoDB, MySQL, PostgreSQL, among others. To efficiently handle and store spatial data, spatial databases rely on a Spatial Database Management System (SDBMS) which extends upon the capabilities of a traditional database management system. SDBMS typically supports three types of queries: *(i)* set operators (e.g., disjoint, touch, contains), *(ii)* spatial analysis (e.g., distance, intersection), and *(iii)* other basic functions such as envelope, boundary, etc. This is efficiently achieved through the reliance on spatial indices (e.g., R-tree, X-tree, GiST). For instance, R-trees [18] represent objects with their minimum bounding rectangle in the next higher level of the tree. The main intuition here is that a query which does not intersect the bounding rectangle also cannot intersect any of the contained objects.

While there are a number of spatial indexing techniques, all techniques require the database server to check a series of coordinate equalities and inequalities in order to determine its spatial index. This clearly poses a problem when dealing with encrypted data objects. For instance, standard encryption of the coordinate system with a semantically secure cryptosystem such as AES, would not preserve any relationship (i.e., to check for equality/inequality) between two points in the coordinate system. While this would be ideal from a cryptographic point of view, it would not be useful for spatial databases, since no "efficient" indexing would be possible on objects.

2.2 System Model

We consider the following system: we assume the existence of mobile nodes $\mathcal{M}$ (e.g., mobile devices, sensors) which "publish" periodic location reports to a Geofencing service that is hosted in the cloud and consists of multiple database servers. In the sequel, we denote by $\mathcal{D}_i$ the i-th database server; for clarity of presentation, we also denote by $\mathcal{D}$ the "logical" database comprising the various database servers. For simplicity and without loss of generality, we assume a 2D bounded area where the nodes can freely move. Each node $M_i \in \mathcal{M}$ periodically sends location beacons to the database servers. These beacons consist of tuples of the form $\langle \mathrm{ID}_i, \mathrm{Loc}_i \rangle$, where ID_i is an identifier of M_i, Loc_i is a vector of the last movement performed by M_i (e.g., Loc could comprise the last and the current coordinates of M_i). We assume that the movement of node M_i is characterized by a velocity distribution $\mathcal{V}_i$, and a path distribution $\mathcal{P}_i$. For simplicity and without loss of generality, we assume that the movement of M_i between two reported coordinates corresponds to a straight line. Current Geofencing applications require indeed that nodes report their last and the current coordinates; by doing so, the Geofencing server can be a stateless server which does not memorize the last coordinate of each node. As we show in this paper, this also serves to increase the privacy of the entire system.

We assume that $\mathcal{D}$ offers location-based services to customers, denoted in the sequel by $\mathcal{S}$. Here, we assume that customers can "subscribe" to events that occur within a specific sub-area of the map for instance, customers want to be notified when users enter/exit their subscribed Geofences.

In the sequel, we denote by S_i the ith subscriber in $\mathcal{S}$; upon receiving a location report from node M_i, $\mathcal{D}$ checks if M_i's reported coordinates are located within a subscribed area. More specifically, $\mathcal{D}$ relies on existing spatial database functionality which can efficiently compute the intersection between a line and polygons. If the movement vector of a node results in non-empty intersection, $\mathcal{D}$ issues a notification message to S_k (e.g., using a URL of S_k stored at $\mathcal{D}$).

2.3 Adversarial Model and Security Requirements

Throughout our analysis, we assume that the nodes are trusted to report their locations correctly. That is, we assume that these devices cannot be compromised by the adversary. Moreover, we assume that the cloud providers (operating the

database servers) are honest-but-curious. More specifically, we assume that each database server will correctly follow the protocol (i.e., authenticate the nodes, output correct notifications) but is interested in acquiring information about the locations of the nodes in the system, and about the queries that are issued by the customers. Ideally, different database servers do not collude; this assumption especially holds when the database servers are hosted by different clouds (e.g., Amazon, Google). Moreover, we assume that the adversary cannot physically track the mobile users to acquire information about their movements. Finally, we assume that the adversary is computationally bounded (i.e., she cannot acquire secrets, break secure encryption functions, etc.).

As mentioned earlier, the main premise behind our work is to design a privacy-preserving solution for a Geofencing service hosted in the cloud, *without* incurring any modifications on the database servers, and while ensuring that $\mathcal{D}$ does not learn any meaningful information about the location of the users and subscriptions in the system. Since the adversary can compromise a database server, we can express these security properties using the following requirements:

Requirement 1—Confidentiality of Stored Records: Each database server should not learn any meaningful information about the stored subscriptions. This can be ensured if the transcript of interaction between $\mathcal{D}$, $\mathcal{M}$, and $\mathcal{S}$ is (computationally) independent of the actual subscription coordinates.

Requirement 2—Confidentiality of Queries: Similar to Requirement (1), each database server should not learn any meaningful information about the location of the nodes. This includes the direction of the movement, the trajectory taken by each node, the distance travelled by each node, etc.

Recall that both Requirements (1) and (2) should be achieved without compromising the functionality of the database server, i.e., while enabling efficient geo-spatial indexing, search over encrypted data, etc. (cf. Section 2.1).

Note that both Requirements (1) and (2) can only ensure confidentiality of the input/output, but do not prevent the possible correlation between the inputs and outputs of the database when subject to location queries by the nodes. This is the case since $\mathcal{D}$ can learn whether a given publish event matches an encrypted subscribe event. In this work, we do not aim at preventing such information leakage. As far as we are aware, the literature features a number of solutions for this problem. These include delaying some queries to ensure k-anonymity [17, 31, 32], relying on bogus queries/subscription to probabilistically provide input/output unlinkability [21], among many others.

3 PrivLoc: Privacy-preserving Outsourced Geofencing Services

In this section, we introduce PrivLoc, our solution which enables the privacy-preserving outsourcing of Geofencing services to the cloud and we thoroughly analyze its security and privacy provisions.

3.1 Overview of PrivLoc

PrivLoc requires that $\mathcal{D}$ only implements the standard Geofencing functionality specified in Section 2.1 without any modification, given inputs from the mobile users. Thus, we see $\mathcal{D}$ as a Geofencing service hosted in the cloud. Nevertheless, PrivLoc ensures that no meaningful information about the location of the devices and subscriptions in the system is leaked to any entity, including $\mathcal{D}$.

In order to achieve these goals, PrivLoc relies on a trusted server $\mathcal{T}$, which mediates the exchange of information between the devices/subscribers on one side, and $\mathcal{D}$ on the other side. More specifically, $\mathcal{T}$ translates both the mobile device locations, and the subscriptions into "scrambled" inputs that are then stored and processed by $\mathcal{D}$. Although the inputs are subsequently hidden from $\mathcal{D}$, PrivLoc ensures that they can be processed using existing geo-spatial indexing algorithms and always result in a correct database lookup. More specifically, the various operations undergone in PrivLoc are:

- Upon receiving each subscription request, $\mathcal{T}$ translates the subscription requested by the subscribers into the appropriate coordinates and stores them at $\mathcal{D}$.
- Upon receiving each sensor location beacon, $\mathcal{T}$ translates the location into the appropriate coordinates and only forwards the transformed beacon to $\mathcal{D}$.

We see $\mathcal{T}$ as an additional service which is run locally to prevent information leakage towards spatial database servers hosted in clouds, such as Google, Amazon, etc. Clearly, for our solution to be effective, the overhead on $\mathcal{T}$ should be minimal. Indeed, in PrivLoc, $\mathcal{T}$ simply has to apply a series of transformations on the received subscriptions and location reports. We stress at this point that the role of $\mathcal{T}$ can be emulated by the mobile nodes and the subscribers themselves, in case these entities pre-share keys. Otherwise, the presence of a trusted service which orchestrates the key management among users and subscribers of the database is required (see [6, 10, 15, 20, 33] for similar assumptions).

PrivLoc introduces two granularity levels when encrypting locations. On the coarse-granular level, PrivLoc relies on a pseudo-random permutation to emulate a strong encryption function, while on the fine-granular level PrivLoc relies on a weaker notion—order-preserving encryption—to preserve the relative location of closely-related points within a sub-area—referred to as *tiles* —so that the spatial operations performed by $\mathcal{D}$ can be applied within each tile without modification. As we describe in Section 3.2, PrivLoc relies on multiple database servers to ensure that any two consecutive location reports can be fitted to at least one tile. The combination of our techniques ensure that location reports pertaining to each tile can be processed using existing functionality of $\mathcal{D}$, without revealing the location or direction of movements. In the following paragraphs, we discuss in details the various translation operations and the load incurred by $\mathcal{T}$ in PrivLoc.

3.2 PrivLoc: Protocol Specification

PrivLoc unfolds as follows. $\mathcal{T}$ first proceeds to dividing the original map by a regular grid into fixed-size square tiles. For privacy reasons, the area covered by

κ	the security parameter, e.g. 128 bit
L	the width and height of a tile
m	the number of tiles in a column, i.e. the height of the map
n	the number of tiles in a row, i.e. the width of the map

Table 1. Global system parameters of PrivLoc.

a tile should fall above a given threshold. In the sequel, we assume the global system parameters shown in Table 1. Furthermore, PrivLoc makes use of the following functions:

- a pseudorandom function PRF : $\mathcal{K} \times \{0,1\}^* \times \ell \rightarrow \{0,1\}^\ell$, which can be realized by a HMAC construction.
- a pseudorandom permutation PRP : $\mathcal{K} \times \{0,1\}^N \rightarrow \{0,1\}^N$, where $N = \lceil \log_2(nm) \rceil$. Exemplary PRP constructions can be found in [8].

In the setup phase, $\mathcal{T}$ generates key k as a uniformly random bit string of a length depending on the intended security level κ.

As mentioned earlier, PrivLoc requires that $\mathcal{T}$ encrypts the coordinates reported by devices before transmitting them to $\mathcal{D}$. Besides hiding location information, one important goal here is to hide the fine-grained user movements (i.e., direction and distance of movement). Recall that, in our setting, the subscribers are interested in knowing whether a device has crossed a Geofence. For the database servers to determine that, it is therefore necessary for them to be able to compare the relative coordinates between the origin and the destination of every sensor movement segment along with the boundaries of Geofences (or subscriptions).

PrivLoc achieves the aforementioned goals by dividing the map into tiles and distorting the original coordinate system into three different variant maps (for the reasoning why, see following paragraphs). By doing so, PrivLoc encrypts the location of each tile within the maps, but ensures that every received location report can be fitted to at least one full tile in one of the maps. Subscriptions that cross more than one tile are also subsequently split to several smaller Geofences that are completely contained in a single tile. Within that tile, PrivLoc further distorts the direction and distance of the user movement while enabling the spatial database to find all intersections between the movement vector and translated Geofences in the distorted tile using existing indexing techniques.

More specifically, PrivLoc applies the encrypt procedure (Algorithm 1); encrypt takes as input real world coordinates (x, y) and transforms them into obfuscated coordinates $(newx, newy)$. encrypt is executed on *(i)* movement vectors to translate the movement end-points, and on *(ii)* subscriptions to translate the south-west, and the north-east coordinates that define the minimum bounding box of each Geofence.

The algorithm encrypt consists of four main routines:

- permuteTiles is used to divide the map into equal-sized tiles and permute these tiles. By doing so, permuteTiles hides the location of the devices within the map.

Algorithm 1 Coordinate Encryption in PrivLoc

Require: Coordinates x, y, map offset z, master key k
Ensure: Encrypted location $newx, newy$

procedure ENCRYPT(x, y, z, k)
  $(Num, dx, dy) \leftarrow$ COORDINATESONTILE(x, y, z)
  $Num \leftarrow PRP(k; n \times m; Num)$
  $(dx, dy) \leftarrow$ ROTATETILE(dx, dy, Num, k)
  $(dx, dy) \leftarrow$ FLIPTILE(dx, dy, Num, k)
  $(dx, dy) \leftarrow \mathrm{OPE}(dx, dy, Num, k)$
  $newx \leftarrow (Num \ \mathbf{mod}\ n) \times L + dx$
  $newy \leftarrow (Num \ \mathbf{div}\ n) \times L + dy$
  return $newx, newy$
end procedure

- rotateTile and flipTile are used to rotate and flip each tile in the distorted map. Both routines serve to hide the direction of devices' movements within the original map.
- OPE is used to hide the distance between any two locations within each tile.

Here, permuteTiles, rotateTile, and flipTile hide the location and direction of the movement of each device from $\mathcal{D}$, while OPE distorts the distances within each tile. All four routines, however, enable $\mathcal{D}$ to rely on existing indexing techniques (c.f. Section 2.1) to compare different locations *within each tile*.

Hiding Movements within a Tile

Recall that PrivLoc operates on movement vectors reported by the users. Upon receiving the vector from devices, $\mathcal{T}$ encrypts the vector, by applying the encrypt procedure in Algorithm 1, to both the start and end-point of the movement vector. These points are represented by 2D coordinates (x, y). The encryption function invokes the procedure coordinatesOnTile, which returns the tile number Num of the tile where the point (x, y) lies on, and the relative coordinates (dx, dy) on the respective tile. As the map is split in $n \times m$ tiles, we have $Num \in \mathbf{Z}_{n \times m}$. This coordinate translation is achieved by means of Algorithm 2.

As the first step of hiding the location of these points within the map, $\mathcal{T}$ permutes the tiles using the key k by applying the pseudorandom function: $PRP(k; n \times m; Num)$. Note that k is only held by the trusted server $\mathcal{T}$. This function computes a pseudorandom permutation of the numbers in $\mathbf{Z}_{n \times m}$ and outputs the new position Num for the tile T.

To further hide the linkability between two consecutive location reports (i.e., to prevent leakage of movement information), $\mathcal{T}$ first rotates each tile using rotateTile (rotation chosen at random among 0, $\pi/2$, π, $3\pi/4$), and then (individually) flips them using flipTile, in order to obfuscate the direction of the movement. Both operations result in a transformation entropy of three bits per tile. When combined with permuteTile, this results in a total of $3\lfloor \log_2(nm) \rfloor$ bits of entropy per tile.

Algorithm 2 Coordinate Translation

Require: Coordinates x, y and offset z of the respective map
Ensure: Num of the tile where x, y lies and relative coordinates dx, dy on the tile

procedure COORDINATESONTILE(x, y)
  $x0 \leftarrow (x - z \textbf{ mod } L \times m) \textbf{ div } L$
  $y0 \leftarrow (y - z \textbf{ mod } L \times n) \textbf{ div } L$
  $Num \leftarrow (x0 + y0 \cdot n)$
  $(dx, dy) \leftarrow x \textbf{ mod } L, y \textbf{ mod } L$
  return$(x0, y0, dx, dy)$
end procedure

Algorithm 3 Tile Flipping

Require: Coordinates dx, dy and Num
Ensure: Coordinates dx, dy after flipping the tile

procedure FLIPTILE(x, y, k)
  $flip \leftarrow \text{PRF}(k;$ "flip", $Num;$ 1)
  if then$flip = '1'$
    $dx' \leftarrow (L - 1) - dx$
  end if
  return(dx', dy)
end procedure

rotateTile is described in Appendix A. rotateTile takes the output of permuteTile; the position of the tile T is not changed in this algorithm. The rotation angle of T is determined using the key-based pseudo-random function PRF.

$\mathcal{T}$ then applies tile flipping (Algorithm 3). Given the outputs of rotateTile, flipTile outputs the newly flipped coordinates $(newx, newy)$. Note that flipTile is analogous to rotateTile, except that the coordinates are transformed by a mirror matrix.

In order to hide the distances of movements executed by devices, PrivLoc further relies on the use of order-preserving encryption (OPE) within each tile. Order-preserving encryption has the property that the relative order among coordinates is preserved after encryption. That is, let $\text{OPE}(p)$ denote the order-preserving encryption of plaintext p. Then, the following holds:

$$\begin{cases} \text{OPE}(p_1) \leq \text{OPE}(p_2) \; if p_1 \leq p_2, \\ \text{OPE}(p_1) > \text{OPE}(p_2) \; if p_1 > p_2. \end{cases}$$

Hiding Cross-Tile Movements

Clearly, encrypt can only hide the trajectory, and distance exhibited by location reports from $\mathcal{D}$ when they fall within the same tile. However, if two consecutive location reports cross the boundary of a single tile and are matched to different tiles, then *(i)* $\mathcal{D}$ cannot compare the relative location advertised by these reports and *(ii)* $\mathcal{D}$ might be able to guess that these tiles correspond to physically connected tiles in the original map.

Thus, a query for a movement which crosses the tile boundary must be avoided in PrivLoc. To achieve that, PrivLoc relies on more than one tiling of the same map (see Figure 1); although a movement might cross the boundary of one map tiling, PrivLoc ensures that there is at least one tiling, in which the movement falls completely within a single tile. Here, by different map tilings, we refer to slightly shifted variants of the original map, which have been processed (i.e., using the encrypt routine) by means of different keys (cf. Figure 1). As shown in Figure 1, the maximum magnitude of a movement vector that PrivLoc accepts is bounded by $l = t/3 + \epsilon$, where t is the length of one tile. Thus, all vectors with magnitude $d(x, y) < r = t/3$ can be located within one complete tile stored at least on one of the three servers. Within that tile, this enables the comparison between all closely located points x and y, where $d(x, y) < t/3$, while ensuring that the movement from location x to y does not cross the tile[1]. An example of a map with three tilings is depicted in Figure 1. As shown in Figure 1, we point out that relying on two tilings of the same map is not enough. This is the case since there exist movement vectors which can still cross two different tiles in any two map tilings. However, it is easy to show that given three shifted tilings of the same map, there is no point where all boundaries intersect, thus ensuring that every movement can be fitted to one tile pertaining to at least one tiling. This is exactly why PrivLoc requires the presence of three database servers $\mathcal{D}_1$, $\mathcal{D}_2$, and $\mathcal{D}_3$, which store subscriptions pertaining to the three map tilings.

All received subscriptions are stored encrypted with three different keys k_1, k_2 and k_3 on all three servers respectively. Whenever $\mathcal{T}$ receives a location report, it will choose to query the database server $\mathcal{D}_i, i \in \{1, 2, 3\}$ for whom this report falls in a full tile (see Appendix B for the algorithm). Here, $\mathcal{T}$ can efficiently find out which server to query for each movement (by essentially performing two integer divisions, cf. Section 4).

As a by-product, we point out that the reliance on multiple database servers in our scheme inherently achieves load balancing of the load on the servers and increases the load capacity of the entire system. This is the case since $\mathcal{T}$ only queries one server for each received location report.

Encrypting Subscriptions

We now proceed to describing how $\mathcal{T}$ encrypts subscriptions and stores them at the database servers $\mathcal{D}$. Clearly, for $\mathcal{D}$ to be able to use existing functionality to compute the intersections of the movements with subscriptions, the subscriptions must be encrypted using the same routine that is used to process the location reports.

More specifically, $\mathcal{T}$ uses the encrypt routine to encrypt the north-east and the south-west coordinates which define the minimum bounding box for each subscription. This is done using the three keys k_1, k_2, and k_3, respectively. The

[1] Thus, this ensures that an adversary who can observe consecutive movements cannot acquire information about the relative position of the actual tiles in each transformed map.

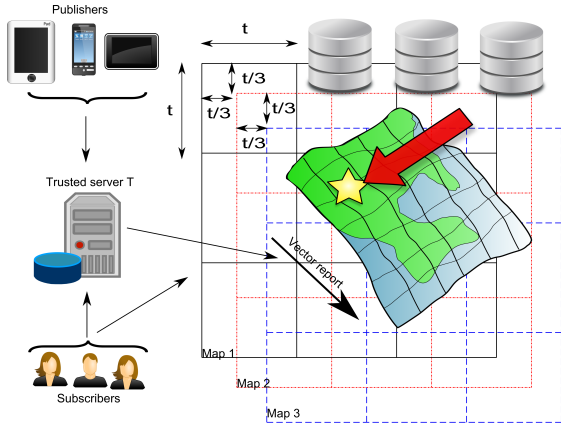

Fig. 1. Example of three tilings of the original map. Notice here that the movement vector crosses the boundaries of tiles in "Map 1" (shown in solid lines) and "Map 2" (shown in dotted lines) but can be fitted complexity within a tile in "Map 3" (shown in dashed lines).

resulting encrypted coordinates are stored in database servers $\mathcal{D}_1$, $\mathcal{D}_2$, and $\mathcal{D}_3$ (see Figure 2). Here, $\mathcal{T}$ must ensure that:

- In case a subscription crosses the boundaries of tiles for the map tiling at a database server, the subscription needs to be split into parts, that each completely fit within one tile of the corresponding map tiling. This process has to be repeated for each of the three servers independently. While this incurs additional storage overhead per server to store the tiles, we show in Section 4 that the storage blowup incurred by PrivLoc can be, to a large extent, tolerated in realistic settings.
- $\mathcal{T}$ batches the upload of $\bar{k}$ subscriptions to the three database servers. Here, $\bar{k}$ denotes a desired privacy threshold. As we show later, this ensures that database servers cannot temporally correlate various subscriptions.

3.3 Security Analysis

Before analyzing the security of PrivLoc, we capture Requirements (1) and (2) (cf. Section 2.3) using the following security game, $\text{Priv}^{\mathcal{A}}$, which involves a p.p.t adversary $\mathcal{A}$ and a challenger $\mathcal{C}$. In our game, $\mathcal{C}$ simulates location reports and subscription requests and emulates the role of the trusted server $\mathcal{T}$ which interacts with $\mathcal{D}$. More specifically, the $\text{Priv}^{\mathcal{A}}$ game unfolds as follows.

Setup. The challenger $\mathcal{C}$ sets up 3 database servers $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$ and generates the respective master keys k_1, k_2 and k_3 using the setup routine.

Run. The challenger $\mathcal{C}$ simulates location reports and subscriptions. More specifically, $\mathcal{C}$ simulates the presence of N nodes with distribution $(\mathcal{V}, \mathcal{P})$, and S

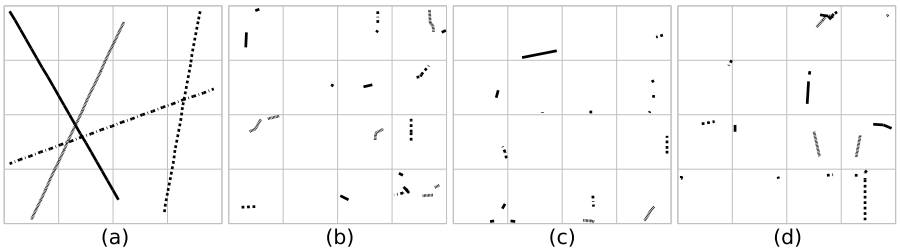

Fig. 2. Exemplary run of PrivLoc with four different movements. Each path displayed in (a) consists of 20 movement reports, each of a length of 30% of the tile width. We depict the snapshots seen by database servers $\mathcal{D}_1$, $\mathcal{D}_2$, and $\mathcal{D}_3$ in (b), (c) and (d), respectively. Note that the marking of the vectors in (b)-(d) is only for visualisation, since different location reports cannot be linked by the database.

subscriptions. On input a location report (x, y), $\mathcal{C}$ executes chooseServer and encrypt(x, y, k).

Compromise. $\mathcal{A}$ chooses *one* server $\mathcal{D}_i^{\mathcal{A}}$, $i \in \{1, 2, 3\}$ at time t. Starting from time t, $\mathcal{A}$ acquires a trace consisting of all inputs and outputs to/from $\mathcal{D}_i^{\mathcal{A}}$.

Challenge. $\mathcal{A}$ then chooses a location report r that arrived at $\tilde{t} > t$ and sends r as a challenge to $\mathcal{C}$.

Response. Upon reception of r, $\mathcal{C}$ locates the node M_k which issued r, and randomly flips a bit b. If $b = 0$, then $\mathcal{C}$ sends to $\mathcal{A}$ a trajectory of movement vectors followed by node M_k. Otherwise, if $b = 1$, $\mathcal{C}$ creates a randomly generated trajectory comprising of movement vectors derived from the distributions $\mathcal{V}_k$ and $\mathcal{P}_k$.

We define the advantage of $\mathcal{A}$ in the above game $\text{Priv}^{\mathcal{A}}$ by:

$$\text{Adv}_{\text{Priv}^{\mathcal{A}}} = \text{Prob}[b' \leftarrow \mathcal{A} : b = b']$$

Definition 1. ***Location Tracking***
We say that a system prevents ϵ-tracking if $\text{Adv}_{\text{Priv}^{\mathcal{A}}} = \frac{1}{2} + \epsilon$.

Clearly, a system perfectly prevents location tracking when ϵ is negligible. Definition 1 captures Requirements (1) and (2) in Section 2.3. Here, $\mathcal{A}$ can compromise one database server. After observing the transcript of interactions with the compromised server, our goal is to prevent $\mathcal{A}$ from acquiring information about the trajectory followed by any node M_k of her choice. This is captured by the fact that the probability that $\mathcal{A}$ can distinguish the trajectory adopted by M_k from any random path is negligible. Notice that $\text{Priv}^{\mathcal{A}}$ captures the ability of the adversary to infer information about the trajectory taken by users using side-information like the velocity of movement, the distance traveled, etc.

Security Analysis: We informally analyze the security of our scheme with respect to the aforementioned $\text{Priv}^{\mathcal{A}}$ game.

Notably, our goal is to show that an adversary $\mathcal{A}$ which compromises one database server $\overline{\mathcal{D}_i}$ and has access to the inputs/outputs of $\overline{\mathcal{D}_i}$ cannot acquire meaningful information about the trajectory of users. Recall that the communication between users, subscribers, and $\mathcal{T}$ is performed over confidential and authenticated channels, which does not give $\mathcal{A}$ any advantage in acquiring information about the (plaintext) location reports and the subscriptions in the system.

In analyzing the advantage of $\mathcal{A}$ in $\text{Priv}^{\mathcal{A}}$, two cases emerge:

1) ***Analyzing single location reports:*** Having received an encrypted location vector from $\mathcal{T}$, we note that $\mathcal{A}$ cannot infer the location of the corresponding user whom generated it in the original map. This is the case since the keyed PRP used in the permuteTiles function ensures that the adversary cannot guess the actual location of the tile which hosts the report in the original map. Moreover, OPE ensures that $\mathcal{A}$ cannot acquire the actual distance travelled by the user while the keyed tile rotation and flipping ensure that $\mathcal{A}$ cannot guess the direction of the user movement given the received location vector.
2) ***Correlating two or more location reports:*** Recall that $\overline{\mathcal{D}_i}$ will only receive location reports which correspond to a complete tile given the tiling hosted by $\overline{\mathcal{D}_i}$. As mentioned earlier, the combined use of the permuted tile location of the report, the OPE-obfuscated distance travelled by the user, and tilted/rotated movement direction does not offer any distinguisher for $\mathcal{A}$ to correlate two or more location reports.

Similarly, it is easy to show that $\mathcal{A}$ cannot acquire any meaningful information about the subscriptions in the system. Note that when tiling and permuting tiles at each server, each subscription might be split into a number of smaller subscriptions within each tile. However, assuming that the system hosts a number of subscriptions, this does not give any advantage to $\mathcal{A}$ in inferring information about the subscriptions. Note, here, that $\mathcal{A}$ can acquire considerable information as $\mathcal{T}$ populates $\overline{\mathcal{D}_i}$. For instance, it is straightforward for $\mathcal{A}$ to guess with high probability that consecutive subscriptions in $\overline{\mathcal{D}_i}$ correspond to the same actual subscription which was subsequently split by $\mathcal{T}$ to prevent subscriptions from crossing tile borders. This is exactly why PrivLoc requires that $\mathcal{T}$ batches the processing of $\bar{k}$ (genuine) subscriptions at a time. On one hand, this enhances the anonymity of the subscriptions, and on the other hand, this prevents $\mathcal{A}$ from correlating stored encrypted subscriptions (e.g., map them to the same subscription) and acquiring information about the tile permutation.

As mentioned in Section 2.3, PrivLoc can only ensure confidentiality of the input/output, and does not prevent the correlation among the inputs and outputs of $\overline{\mathcal{D}_i}$. That is, $\mathcal{A}$ can learn whether a given location report has triggered a response from $\overline{\mathcal{D}_i}$ and therefore corresponds to a user entering/exiting a Geofence. The literature features a number of solutions for this problem (see [17,21,31,32]).

4 Implementation & Evaluation

In this section, we implement PrivLoc and we evaluate the performance of our implementation in comparison to the setting where the Geofencing service is not outsourced to the cloud, but is locally offered using an existing spatial database.

4.1 Implementation Setup

We implement the PrivLoc service using Java. In our implementation, we set $\kappa = 128$, we adapt the OPE from [9], and we use HMAC-SHA1 as the PRF in the rotateTile and flipTile routines. The tile permutation is achieved by an array permutation given a secure random number generator. In our setup, we deploy PrivLoc on an 8-core Intel Xeon E3-1230 with 16 GB RAM; the various clients issuing location reports and/or subscriptions were co-located with $\mathcal{T}$ on the same machine.

We chose a square $100km \times 100km$ map in which the maximum size of a subscription is $100m \times 100m$ (this was also the tile size in our implementation). Conforming with Section 3.2, we set the maximum movement distance of a user between two consecutive location reports to $\frac{t}{3} \approx 30m$. We rely on the random waypoint model to simulate node mobility. More specifically, for each user in the system, we assume a random initial location in the map; for each subsequent movement, the movement vector angle is chosen randomly from [0,360] degrees, while the distance of each movement is chosen randomly from [0,30] meters. Throughout our evaluation, we assume a realistic setting where the ratio of location reports to incoming subscription requests is 19:1.[2]

Based on this setup, we measure the throughput achieved by PrivLoc with respect to the average latency incurred on $\mathcal{T}$ due to a location report or a subscription. To increase the load on $\mathcal{T}$, we vary the number of concurrent clients from 1 to a maximum of 256. In our experiments, we are interested in assessing the performance of PrivLoc when the links between $\mathcal{T}$ and the database servers are not the bottleneck. For that purpose, we short-circuit the database servers and we abstract away the time to upload the location reports by $\mathcal{T}$ to the database servers. Under these settings, we compare the throughput and latency achieved by PrivLoc to the traditional case where the spatial database is entirely hosted locally; here, we compare the performance of PrivLoc to *(i)* a stand-alone MySQL spatial database (with R-tree indexing) and *(ii)* a PostgreSQL database (with GiST indexing). Both databases were deployed within our 8-core Intel Xeon E3-1230 with 16 GB RAM; we report their performance with different subscription table sizes (i.e., initial number of records). We also evaluated PrivLoc when compared to MongoDB. Since the performance exhibited MongoDB was far inferior to that of MySQL and PostgreSQL, we omit these measurements from our evaluation.

[2] We conducted experiments where the ration between where location reports and subscriptions is 1:1. Our results were similar to the setting featuring a ratio of 19:1.

4.2 Evaluation Results

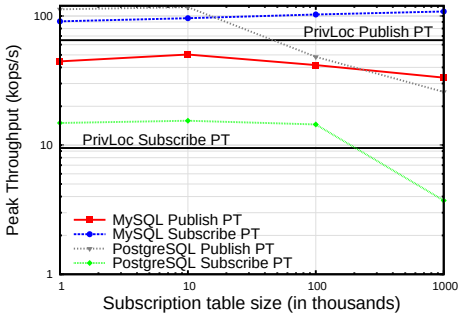

Fig. 3. Comparison between PrivLoc and local MySQL and PostgreSQL databases.

We start by evaluating the performance of PrivLoc w.r.t. to the MySQL and PostgreSQL databases while varying the initial number of stored subscriptions. This experiment captures the performance of PrivLoc compared to locally hosted spatial databases as the number of subscription populating these databases increases with time. For that purpose, we measure the peak throughout (PT) exhibited by PrivLoc and compare it with that of MySQL, and PostgreSQL respectively, in settings where the initial number of stored subscriptions varies from 1,000 to 1,000,000. Our results are depicted in Figure 3. Our findings show that the PT of PrivLoc is superior to MySQL when processing location reports, irrespective of the number of stored subscriptions. Here, PrivLoc outperforms PostgreSQL as the number of subscriptions stored by the PostgreSQL database increases beyond 50,000 records. This is also the case when processing subscription requests. Nevertheless, our results indicate that, even when the initial number of subscriptions stored in the MySQL, and PostgreSQL databases is as low as 1,000, the relative PT achieved by PrivLoc can be easily tolerated.

In a second experiment, we evaluate the relative performance of PrivLoc in the realistic case where the Geofencing service has been running for some time, and has accumulated the subscriptions from a large number of subscribers. To simulate this case, we insert 1,000,000 subscription records in the databases. Figure 4 depicts the latency incurred in PrivLoc for the processing of location reports and subscriptions in the system with respect to the achieved throughput (measured in the number of 1000 operations per second). For comparison purposes, we also include the performance achieved by a local MySQL and a PostgreSQL spatial database in the same setup. Our results (cf. Figure 4(a)) show that PrivLoc is at least twice faster than locally processing location reports on both spatial databases. Moreover, the peak throughput achieved in PrivLoc is more than 2 times higher than that achieved by the MySQL and PostgreSQL databases. However, our results in Figure 4(b) show that the latency and peak throughput achieved by PrivLoc are modest when inserting subscriptions, compared to the local MySQL database. This is due to the blow-up in the number of subscriptions. Notably, since $\mathcal{T}$ splits the map into small tiles, a subscribed area might be further split by $\mathcal{T}$ if it crosses multiple tiles. Note that each subscribed area results in an average of 6.75 subscribed areas which will be encrypted by PrivLoc and pushed to the $n = 3$ database servers. This blow-up is dependent on the number of database servers, and the size of the subscribed area with respect to the size of the tiles. More specifically, the average area of a subscription in our setup is 1/4 of the area of a tile; this means that the average number of partitions

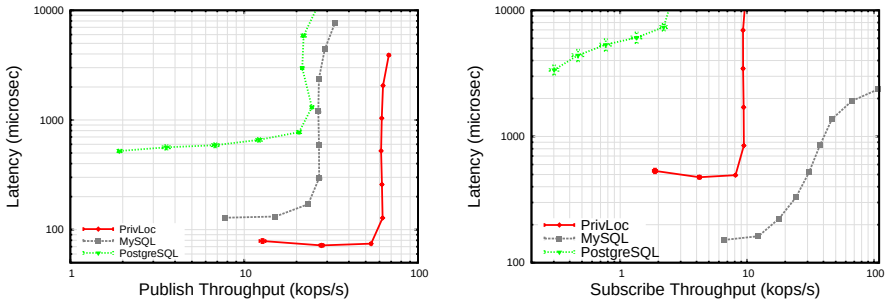

(a) Average latency for processing a location report by $\mathcal{T}$ with respect to the throughput.

(b) Average latency for processing a subscription by $\mathcal{T}$ with respect to the throughput.

Fig. 4. Throughput and latency in PrivLoc. Here, we assume a subscription table size of 1,000,0000 records. Each data point in our measurements is averaged over 10 independent measurements; we present the corresponding 95% confidence intervals.

of a subscribed area on each database server is $N_p = 1 \cdot \frac{1}{4} + 2 \cdot \frac{2}{4} + 4 \cdot \frac{1}{4} = 2.25$. The total blowup in terms of subscriptions is subsequently $3N_p = 6.75$.

In spite of the storage blowup, our findings nevertheless show that PrivLoc considerably outperforms a local PostgreSQL database.

5 Related Work

In what follows, we briefly overview existing contributions in the area. In [7], Barkkuus and Dey show that users were concerned about the ability of services to track them.

Most privacy-enhancing solutions for location-based services rely on a trusted "location anonymizer" service which hides the location of users. These services either provide k-anonymity [13, 17, 26, 31, 32] or spatial/temporal cloaking with an area of interest [6, 10, 15, 20, 33]. A number of solutions rely on inserting fake queries in order to prevent a database server from learning the actual location reports (e.g., [21]). While these solutions provide k-anonymity, they incur significant additional costs on the database server. Other solutions rely on location perturbation/obfuscation; these solutions map the location reports to a set of pre-defined landmarks [19] or blur the user location into a spatial area using linear transformations [6, 15, 20, 22, 34] Such solutions indeed hide the location of users but might affect the accuracy of the location-based service. Moreover, these solutions can only hide the location of a user, but do not aim at hiding the user movement.

To prevent location tracking, Gruteser and Liu [16] propose disclosure control algorithms which hide users' positions in sensitive areas and withhold path information that indicates which areas they have visited. Other schemes rely on Private Information Retrieval (PIR) algorithms in order to enable privacy-preserving queries in spatial databases [14, 24]. PIR schemes allow a querier to

retrieve information from a database server without revealing what is actually being retrieved from the server. However, these solutions are computationally intensive and require modifications to the database server in order to process the blurred location queries.

In [25], Pfitzmann *et al.* define unlinkability and privacy in pseudonymous systems. Dwork [12] define differential privacy and quantify information leakage from the query access of individuals. In [30], Shokri *et al.* quantify location privacy using the error of the adversarial estimate from the ground truth. In [11, 27], various entropy-based metrics are introduced to assess the communication privacy in anonymous networks.

6 Conclusion

In this paper, we proposed a novel solution, PrivLoc, which enables privacy-preserving outsourcing of location-based services to the cloud without leaking any meaningful information to the cloud provider. PrivLoc goes one step beyond existing solutions in the area and targets Geofencing services where users send a vector of their movements for the service provider to detect whether a user has crossed a given Geofence. We analyze the security and provisions of PrivLoc and show that PrivLoc does not leak information about the location, movement, trajectory, and/or velocity of the users to the Geofencing database. Our evaluation of PrivLoc shows that the overhead incurred by our solution can be largely tolerated in realistic deployment settings.

Acknowledgements

This work was supported by the EU FP7 SMARTIE project (contract no. 609062).

References

1. Foursquare. Available from `http://foursquare.com/`.
2. Keep Track of Your Kids With Geofences. Available from `http://netsecurity.about.com/od/newsandeditoria2/a/Keeping-Tabs-On-Your-Kids-With-Geofences.htm`.
3. Location Based Notifications. Available from `http://www.plotprojects.com/`.
4. Valtus Spatial Data Cloud Services. Available from `http://www.valtus.com/products/vault-spatial-data-cloud-services/`.
5. Yelp. Available from `http://www.yelp.com/`.
6. Bhuvan Bamba, Ling Liu, Peter Pesti, and Ting Wang. Supporting Anonymous Location Queries in Mobile Environments with Privacygrid. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, pages 237–246, New York, NY, USA, 2008. ACM.
7. Louise Barkhuus and Anind Dey. Location-based services for mobile telephony: a study of users privacy concerns. In *INTEREACT*, pages 709–712, 2003.
8. John Black and Phillip Rogaway. Ciphers with Arbitrary Finite Domains. In *CT-RSA 2002*, volume 2271 of *Lecture Notes in Computer Science*, pages 114–130. Springer, 2002.

9. Alexandra Boldyreva, Nathan Chenette, Younho Lee, and Adam O'Neill. Order-preserving symmetric encryption. In *Proceedings of the 28th Annual International Conference on Advances in Cryptology: The Theory and Applications of Cryptographic Techniques*, EUROCRYPT '09, pages 224–241, Berlin, Heidelberg, 2009. Springer-Verlag.
10. Chi-Yin Chow and Mohamed F. Mokbel. Enabling Private Continuous Queries for Revealed User Locations. In *Proceedings of the 10th International Conference on Advances in Spatial and Temporal Databases*, SSTD'07, pages 258–273, Berlin, Heidelberg, 2007. Springer-Verlag.
11. C. Diaz, S. Seys, J. Claessens, and B. Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, April 2002.
12. C. Dwork. Differential privacy: a survey of results. In *Proceedings of the 5th international conference on Theory and applications of models of computation*, TAMC'08, 2008.
13. Buğra Gedik and Ling Liu. Protecting Location Privacy with Personalized k-Anonymity: Architecture and Algorithms. *IEEE Transactions on Mobile Computing*, 7(1):1–18, January 2008.
14. Gabriel Ghinita, Panos Kalnis, Ali Khoshgozaran, Cyrus Shahabi, and Kian-Lee Tan. Private Queries in Location Based Services: Anonymizers Are Not Necessary. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 121–132, New York, NY, USA, 2008. ACM.
15. Marco Gruteser and Dirk Grunwald. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In *Proceedings of the 1st International Conference on Mobile Systems, Applications and Services*, MobiSys '03, pages 31–42, New York, NY, USA, 2003. ACM.
16. Marco Gruteser and Xuan Liu. Protecting Privacy in Continuous Location-Tracking Applications. *IEEE Security and Privacy*, 2(2):28–34, March 2004.
17. Saikat Guha, Mudit Jain, and Venkata N. Padmanabhan. Koi: A Location-privacy Platform for Smartphone Apps. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, NSDI'12, pages 14–14, Berkeley, CA, USA, 2012. USENIX Association.
18. Antonin Guttman. R-trees: A Dynamic Index Structure for Spatial Searching. In *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, SIGMOD '84, pages 47–57, New York, NY, USA, 1984. ACM.
19. Jason I. Hong and James A. Landay. An Architecture for Privacy-sensitive Ubiquitous Computing. In *Proceedings of the 2Nd International Conference on Mobile Systems, Applications, and Services*, MobiSys '04, pages 177–189, New York, NY, USA, 2004. ACM.
20. Panos Kalnis, Gabriel Ghinita, Kyriakos Mouratidis, and Dimitris Papadias. Preventing Location-Based Identity Inference in Anonymous Spatial Queries. *IEEE Trans. on Knowl. and Data Eng.*, 19(12):1719–1733, December 2007.
21. Hidetoshi Kido, Yutaka, Yanagisawa, and Tetsuji Satoh. An Anonymous Communication Technique using Dummies for Location-based Services. In *ICPS*, pages 88–97, 2005.
22. Mohamed F. Mokbel, Chi-Yin Chow, and Walid G. Aref. The New Casper: Query Processing for Location Services Without Compromising Privacy. In *Proceedings of the 32Nd International Conference on Very Large Data Bases*, VLDB '06, pages 763–774. VLDB Endowment, 2006.

23. Jussi Myllymaki and James Kaufman. High-performance spatial indexing for location-based services. In *Proceedings of the 12th International Conference on World Wide Web*, WWW '03, pages 112–117, New York, NY, USA, 2003. ACM.
24. Femi Olumofin, Piotr K. Tysowski, Ian Goldberg, and Urs Hengartner. Achieving Efficient Query Privacy for Location Based Services. In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies*, PETS'10, pages 93–110, Berlin, Heidelberg, 2010. Springer-Verlag.
25. Andreas Pfitzmann and Marit Hansen. Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management ? A Consolidated Proposal for Terminology. *Fachterminologie Datenschutz und Datensicherheit*, pages 111–144, 2008.
26. P. Samarati. Protecting Respondents' Identities in Microdata Release. *IEEE Trans. on Knowl. and Data Eng.*, 13(6):1010–1027, November 2001.
27. A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
28. Shashi Shekhar, Sanjay Chawla, Siva Ravada, Andrew Fetterer, Xuan Liu, and Chang-tien Lu. Spatial databases-accomplishments and research needs. *IEEE Trans. on Knowl. and Data Eng.*, 11(1):45–55, January 1999.
29. Anmol Sheth, Srinivasan Seshan, and David Wetherall. Geo-fencing: Confining Wi-Fi Coverage to Physical Boundaries. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, pages 274–290, Berlin, Heidelberg, 2009. Springer-Verlag.
30. R. Shokri, G. Theodorakopoulos, J. Le Boudec, and J. P. Hubaux. Quantifying location privacy. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2011.
31. Latanya Sweeney. Achieving K-anonymity Privacy Protection Using Generalization and Suppression. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):571–588, October 2002.
32. Latanya Sweeney. K-anonymity: A Model for Protecting Privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, October 2002.
33. Yu Wang, Dingbang Xu, Xiao He, Chao Zhang, Fan Li, and Bin Xu. L2P2: Location-aware location privacy protection for location-based services. In *INFOCOM*, pages 1996–2004, 2012.
34. Man Lung Yiu, Gabriel Ghinita, Christian S. Jensen, and Panos Kalnis. Enabling Search Services on Outsourced Private Spatial Data. *The VLDB Journal*, 19(3):363–384, June 2010.

A Tile Rotation in PrivLoc

Require: Coordinates dx, dy and Num
Ensure: Coordinates dx, dy after rotating the tile

procedure ROTATETILE(x, y, k)
  rotation ← PRF(k; "rotate", Num; 2)
  if rotation = '00' **then**
    $dx' \leftarrow dx$
    $dy' \leftarrow dy$
  else if $rotation = '01'$ **then**
    $dx' \leftarrow dy$

```
        dy' ← (L − 1) − dx
    else if rotation = '10' then
        dx' ← (L − 1) − dx
        dy' ← (L − 1) − dy
    else if rotation = '11' then
        dx' ← (L − 1) − dy
        dy' ← dx
    end if
    return (dx', dy')
end procedure
```

B Querying the appropriate database server in PrivLoc

Require: Coordinates $x0, y0$ and $x1, y1$ describing a movement
Ensure: Server where $x0, y0$ and $x1, y1$ are on the same tile or $\perp$

```
procedure chooseServer(x, y)
    for all D_i, z_i with i = 1, 2, 3 do          ▷ Server D with tile offset z
        nx ← (x0 − z_i mod L × m) div L
        ny ← (y0 − z_i mod L × n) div L
        Num0 ← (nx + ny · n)
        nx ← (x1 − z_i mod L × m) div L
        ny ← (y1 − z_i mod L × n) div L
        Num1 ← (nx + ny · n)
        if Num0 = Num1 then
            return i
        end if
    end for
    return ⊥
end procedure
```